

ASC HPEMS Xyce™ Circuit Simulator: Linear Solver Technology

June 23, 2004

R. Hoekstra

D. Day, M. Heroux, E. Keiter, S. Hutchinson,
T. Russo, E. Rankin, R. Pawlowski

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,
for the United States Department of Energy under contract DE-AC04-94AL85000.

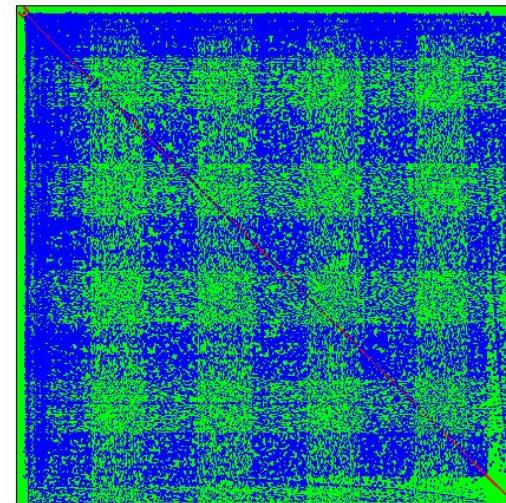


Overview

- **What's Hard about Circuit Problems?**
- **Trilinos Solver Library**
- **Linear Solution for Circuits**
 - Direct
 - Iterative
 - Singleton Filtering
 - Ordering/Partitioning
 - Block Triangular Factorization
- **Conclusions**

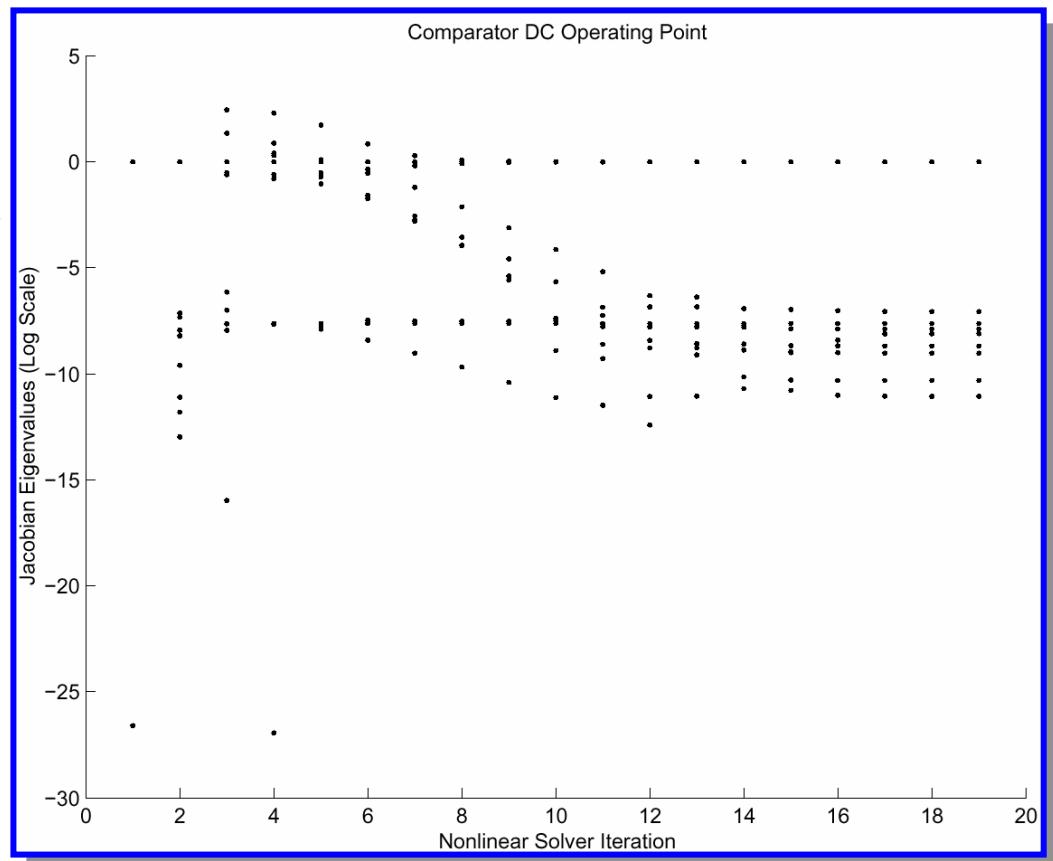
What's HARD?

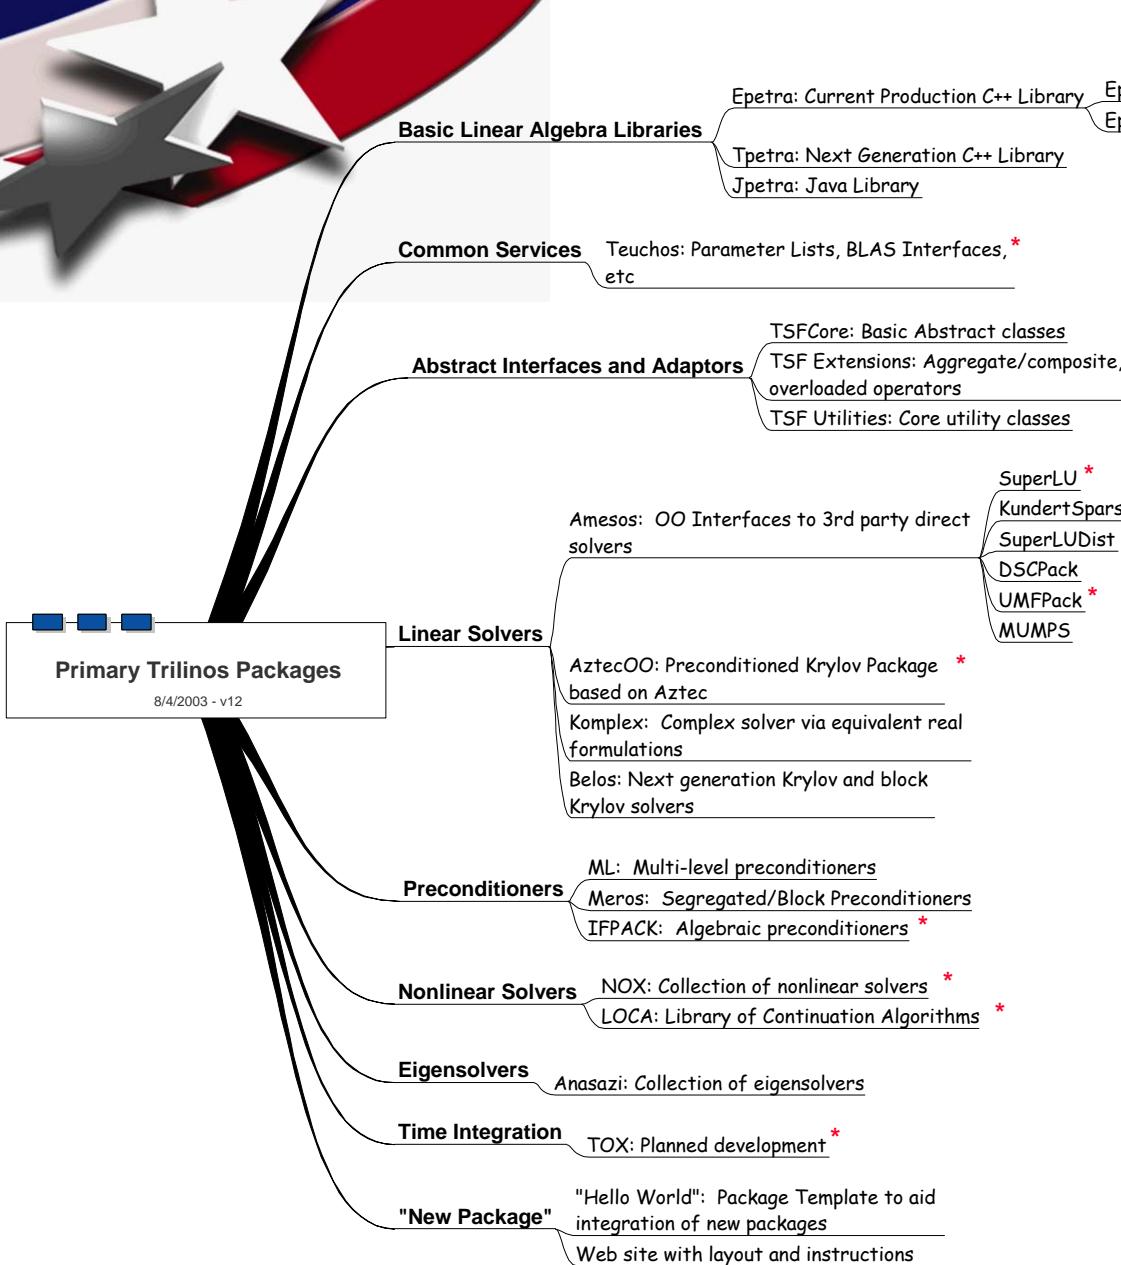
- Stiff coupled DAE's
- Highly Nonlinear Devices
 - Discontinuities, Hysteresis
- Sparse Linear Systems
 - III-Conditioned/Scaled
 - Non-Symmetric
 - Network Topology (Not A Mesh!) Preconditioning
 - Dense Rows Partitioning/
Ordering



Eigenvalues

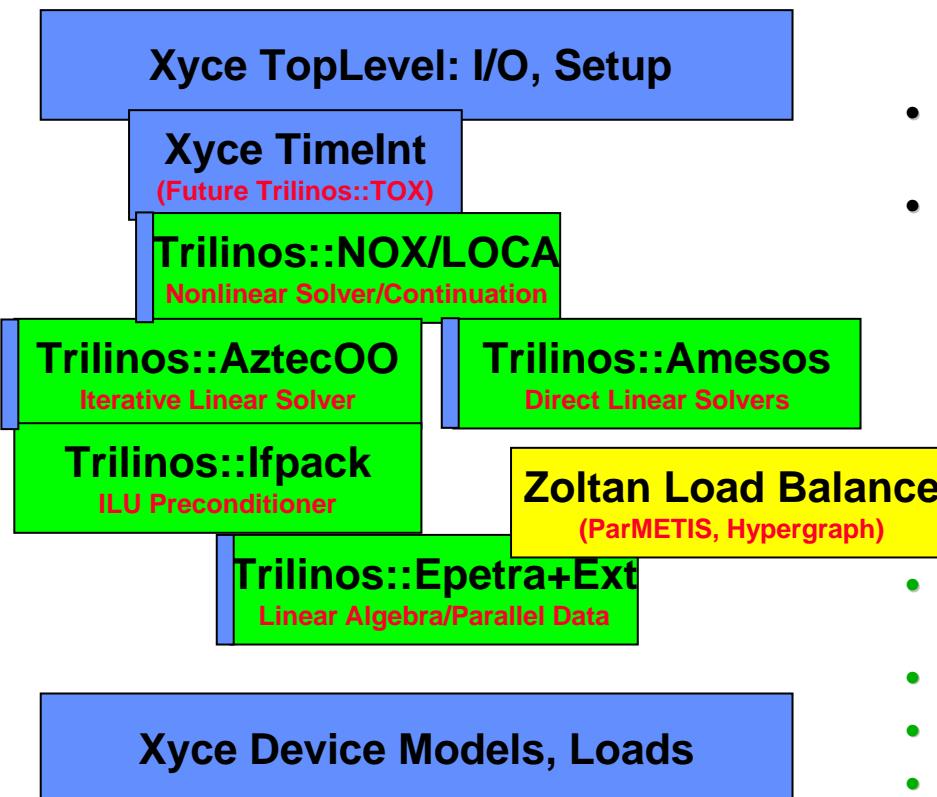
- ⊕ *Newton's Method*
- ⊕ *Spectrum of Jacobian each iteration*
- ⊕ *Semi-Log Plot*
- ⊕ *Numerically Singular until Convergence*
- ⊕ *Non-Uniqueness*





- **Trilinos is a collection of *Packages*.**
- **Focused Package Development**
 - State-of-the-art algorithms in a given problem regime.
 - Small development team of domain experts.
 - Self-contained
 - Individual Configure/Build/Documentation
- **Benefits**
 - Common Infrastructure
 - Common Tools
 - Interoperability
- <http://software.sandia.gov/trilinos>

Xyce(Trilinos/Zoltan)



- Open source libraries under rapid development
- Benefits
 - State-of-the-Art Algorithms
 - Rapid Support
 - Gnu Autotool Configure/Build Environment
- NOX/LOCA: Globalized Newton Type Methods, Homotopy/Continuation
- AztecOO: Preconditioned GMRES
- Ifpack: Enhanced Block-ILUK
- Epetra(Ext): Distributed Memory Linear Algebra and Transformations
- Zoltan: Parallel Partition/Load Balance

Amesos: Sparse Direct Linear Solvers

Library Name	Lang.	Comm.	# procs for solution ¹
KLU	C	Serial	1
UMFPACK	C	Serial	1
Mod. KundertSparse	C	Serial	1
SuperLU_DIST 2.0	C	MPI	any
MUMPS 4.3.1	F90	MPI	any
ScaLAPACK	F77	MPI	any

Iterative Linear Solve

- Trilinos: Epetra, IfPack, AztecOO (**Belos, TSF**)
- Strategy : GMRES
 - Domain Decomposition
 - » Singleton Filtering → **DENSE ROWS**
 - » Zoltan/ParMETIS Partitioning
 - » Overlapping Additive Schwarz
 - AMD/RCM Block Reordering
 - Row/Col Scaling
 - Stabilized ILUT(B)
 - » $B :=$ dual threshold *a priori* diagonal perturbation (A)

Adaptive ILU Preconditioning

- Idea:** Compute ILU factor of a matrix B that is “nearby” original matrix A , but better conditioned. (Generalization of Manteuffel shift)
- Sets up a continuum of preconditioners between accurate but poorly conditioned ILU factor and Jacobi scaling.**
- B differs from A only on diagonal**
- Adaptive Algorithm to test threshold values**

$$b_{ii} = \text{sgn}(a_{ii})\alpha + (1 + \rho)a_{ii}$$

$$\begin{aligned}|b_{ii}| &\geq \alpha \\ |b_{ii}| &\geq (1 + \rho) |a_{ii}|\end{aligned}$$

Max Condest	Iters (Res)	Result	Action	Time
inf	None	Condest not < Condest thresh	Set $\alpha = 10^{-12}$	
10^{12}	1 (10^0)	Hessenberg III-conditioned	Set $\rho = 10^{-2}$	
10^2	10 (10^{-7})	Converged		0.35

Singleton Filtering

- **Row Singleton**
 - Pre-Process

$$\begin{bmatrix} & a_{1j} & & \begin{bmatrix} x_1 \\ \vdots \\ x_j \\ \vdots \\ x_n \end{bmatrix} & = & \begin{bmatrix} b_1 \\ \vdots \\ b_i \\ \vdots \\ b_n \end{bmatrix} \Rightarrow x_j = b_i / a_{ij} \\ & a_{2j} & & & & \\ & \vdots & & & & \\ & \vdots & & & & \\ 0 & \cdots & 0 & a_{ij} & 0 & \cdots & 0 \\ & & & \vdots & & & \vdots \\ & & & a_{nj} & & & \vdots \\ & & & & & & \end{bmatrix}$$

- **Col Singleton**
 - Post-Process

$$\begin{bmatrix} & 0 & & \begin{bmatrix} x_1 \\ \vdots \\ x_j \\ \vdots \\ x_n \end{bmatrix} & = & \begin{bmatrix} b_1 \\ \vdots \\ b_i \\ \vdots \\ b_n \end{bmatrix} \Rightarrow x_j = (b_i - \sum_{k \neq j} a_{ik}x_k) / a_{ij} \\ & 0 & & & & \\ & \vdots & & & & \\ & 0 & & & & \\ a_{i1} & \cdots & \cdots & a_{ij} & \cdots & \cdots & a_{in} \\ & & & 0 & & & \vdots \\ & & & 0 & & & \vdots \end{bmatrix}$$

Dist. Memory Algorithm in Trilinos::EpetraExt

Singleton Filtering

Number of...	Original Problem	After Filtering
Successful Steps Taken	925	892
Failed Steps Attempted	67	53
Jacobians Evaluated	2488	2399
Linear Solves	2488	2399
Failed Linear Solves	32	13
Linear Solver Iterations	33595	12152

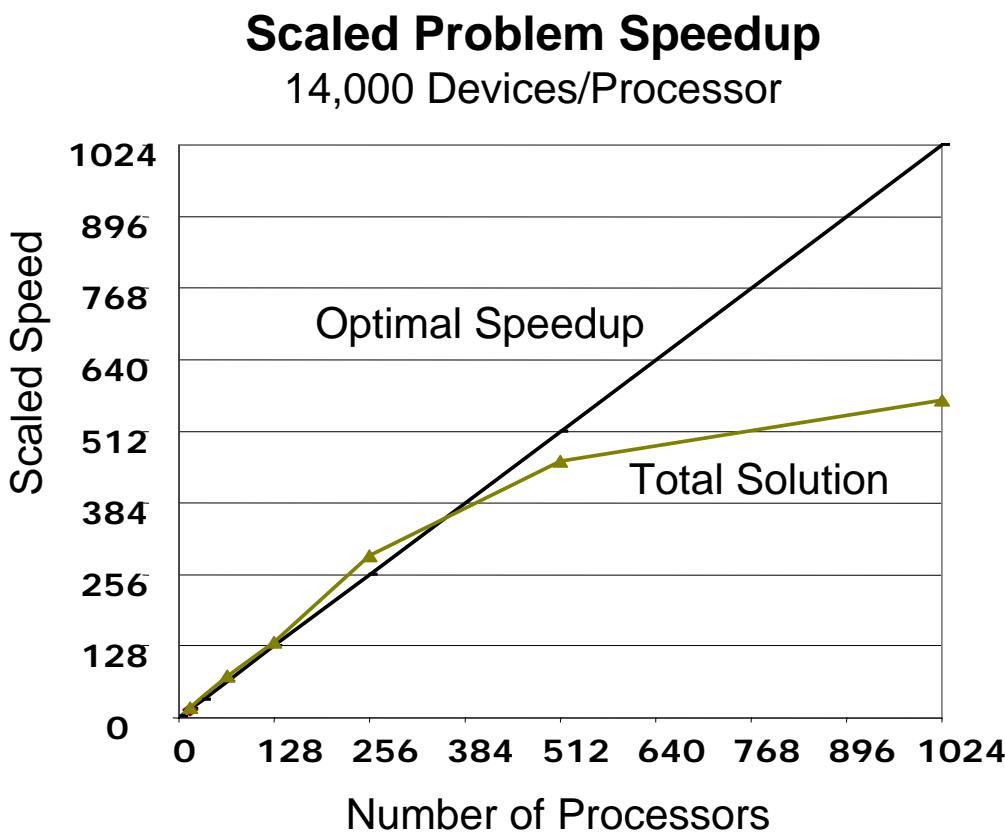
Putting It Together

- Digital Adder on 8 processors
- Improved both Scalability & Robustness
- Maybe Iterative Solvers are viable for Circuits!



	N	Total Cuts	CondEst	GMRES Iters	LinSolve Time	Newton Steps
PC	1220	~1000	3.00E+05	500	4.7	72
PC+SF+PL+RCM+SCALE	1220	188	2.00E+05	200	1.1	45
	1054	68	1.00E+04	127	0.43	53

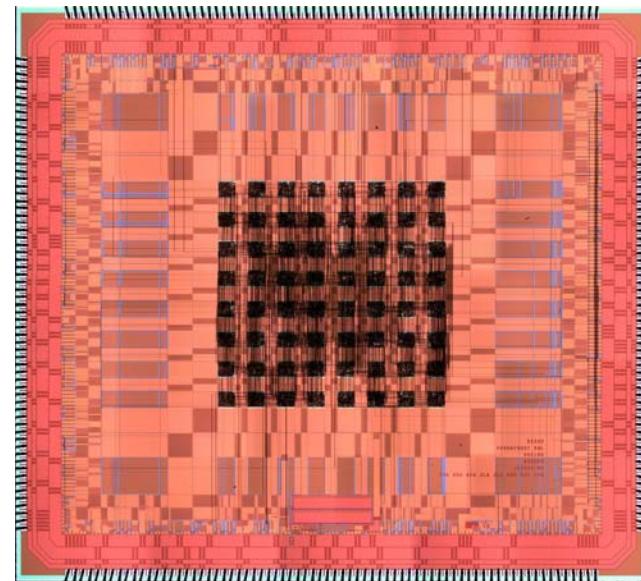
Parallel Scaling



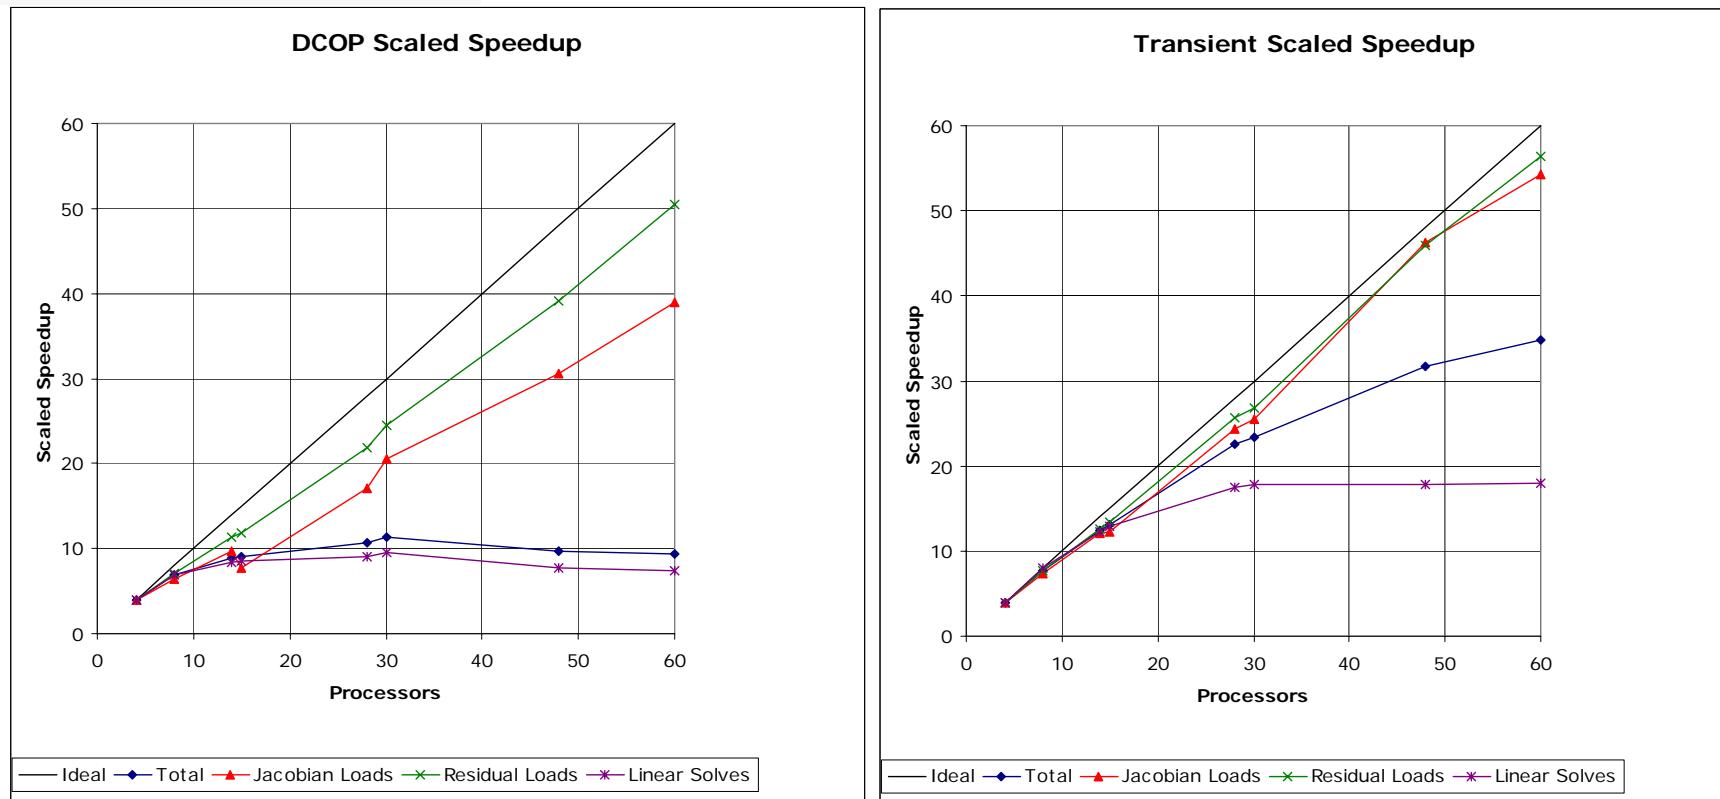
- Nonlinear transmission line
- 14 million devices
- 6 million Unk's
- Over factor of 500 speedup using 1024 processors of ASC White
 - 14,000 electrical devices per processor

Sandia ASIC Design

- **Sandia ASIC Design**
 - Digital circuit. ~250K Transistors.
- **Problem Setup**
 - Distributed memory scalability
- **Init. Cond. : ROBUSTNESS!**
 - Homotopy (NOX/LOCA)
 - Singleton Filtering (EpetraExt)
 - AMD Ordering (EpetraExt)
 - Mod. BILUK Precond. (Ifpack)
- **Transient : PERFORMANCE!**
 - Dominates Run Time!
 - Communication Enhancements (Epetra)
 - Zoltan Partitioning (EpetraExt)

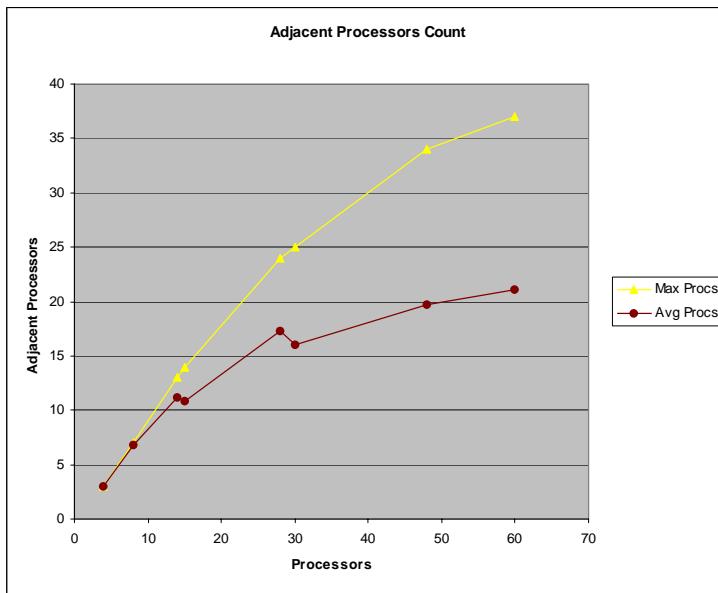
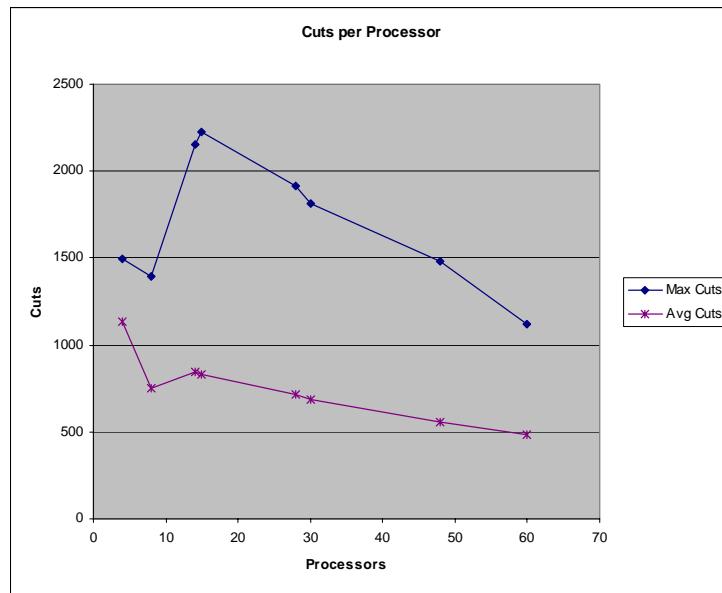


Xyce™ Parallel Scaling Fixed Size ASIC Problem



- Linear solver convergence dominates scalability for DCOP.
- Transient simulation dominates overall runtime; scalable to 32 processors.
- Scaling rolloff corresponds to ~8k devices and ~3k unknowns per processor.

Partitioning Issues



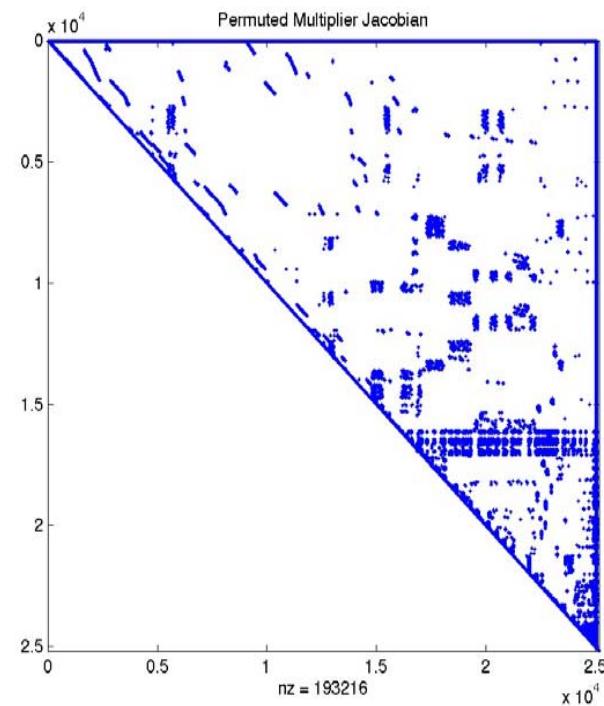
- **Scalable communication volume (cuts)**
- **Not so scalable communication count (adj procs)**
 - Hierarchical nature of circuits?
 - Will comm. count plateau for bigger problems?

Load Balance/Partitioning

- Good but not great success so far...
- New Ideas
 - Weighted Graph Partitioning
 - » Improve BILU Preconditioning but keeping fill in block
 - Reduce max values of off block diagonals by several orders of magnitude for some problems
 - Multi-Constraint Partitioning
 - » Balance Load(Circuit) and Solve(LinSys) Partitions
 - Hypergraph Partitioning
 - » Better representation of non-symmetric systems
 - » Better representation of MatVec communication
 - » Demonstrated as much as 50% communication volume reduction for sample Xyce problems

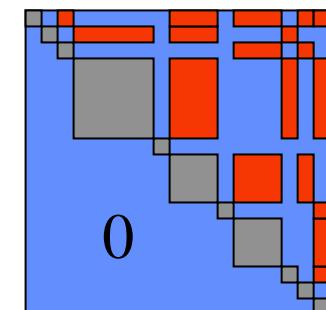
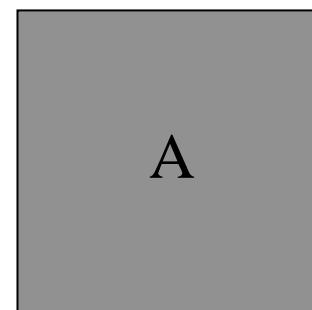
Block Triangular Factorization

- Steady State Analog Circuit Problem are Block “REDUCIBLE”!
- Largest Blocks found: <150
- Novel Algorithm: $O(n_b \cdot s_b^3 + n_b^2 \cdot s_b^2)$
- Current implementation beats our fastest sparse direct solver for $n > 10,000$
- Ill-conditioned ($> 10^{16}$) diagonal blocks can be better managed.

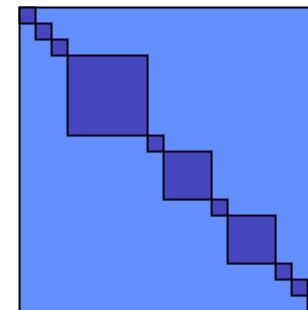


Block Triangular Solve

Block Triangular Factorization
(Alex Pothien's Algorithm)



Invert Diagonals (e.g. SVD, LU)



Block Backsolve

Singular Value Thresholding

- Managing III-Conditioned Diagonal Blocks
- Abs/Rel Thresholding
- Relative to Nonlinear Norm

$$th_i = \max\left(th_{rel}s_{i,\max} + th_{abs}, th_{rel_norm}\|r\|_\infty\right)$$

N	Num Blocks	Largest Block	Cond Est	Ksparse Newton Steps	BTS Newton Steps	Ksparse Solve time	BTS Solve Time	Ksparse Max Norm	BTS Max Norm
16	13	2	2.00E+10	20	12	3.00E-05	4.00E-04	5.00E+00	5.00E+00
1220	188	101	5.00E+15	50	45	7.00E-04	2.00E-01	1.00E+28	1.00E+04
1815	471	109	4.00E+12	61	55	2.00E-03	1.00E-01	2.00E+04	2.00E+04
25187	7401	79	1.00E+23	FAILED	255	11.5	3.5	1.00E+48	2.00E+12

BTS: What Else?

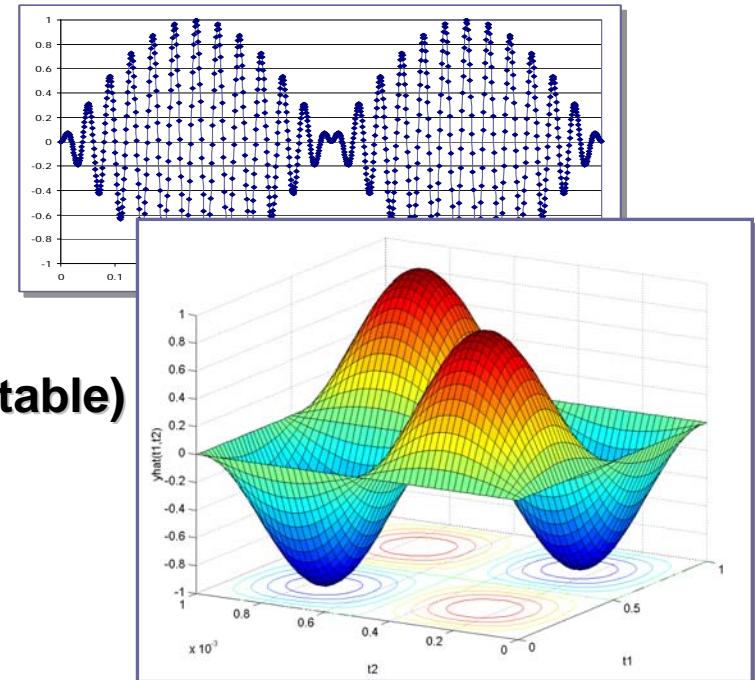
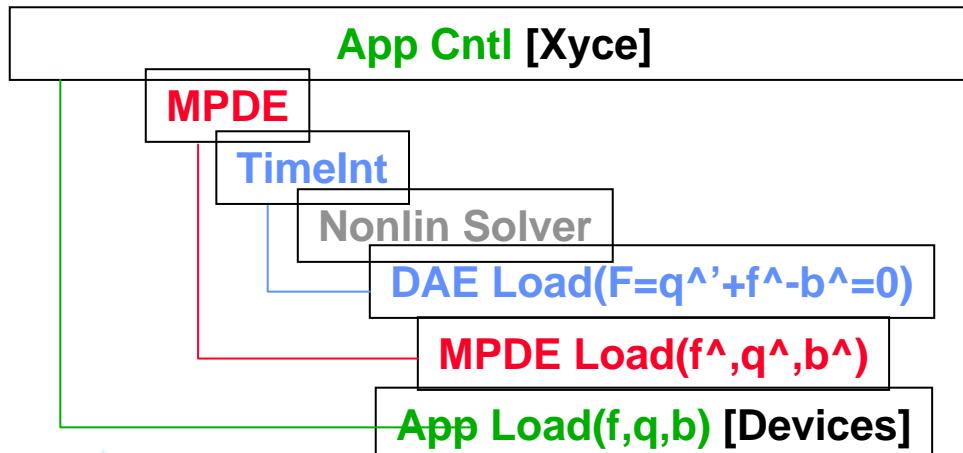
- **Parallel Algorithm**
 - BTF Reordering
 - Invert
 - Backsolve
- **Diagonal Block Inversion**
 - Performance
 - Ill-Conditioning Management
- **Iterative Solver Preconditioning**
- **Nonlinear Algorithm: Step through the diagonal block nonlinear problems**

Future Directions

- **Preconditioned Iterative Solvers**
 - Multi-Level Preconditioners: pARMs, etc.
 - BTF based Preconditioner
 - Intelligent Partitioning for Preconditioning
- **Block Triangular Form**
 - Parallel
 - Managing Ill-Conditioning
 - Direct Solvers (KLU, T. Davis)
- **Partitioning/Load Balance**
 - HyperGraph
 - Multi-Constraint

Time-Parallel Multi-time PDEs

- Beta Capability in Xyce
 - Primary Infrastructure
 - » To be refactored as Trilinos Pkg
 - Block Linear Algebra Manipulation
 - Fast Time Scale Discretization
 - » Arbitrary Order BD and CD (FD unstable)
 - » Freq. Domain to be added



$$x(t) = \sin(2\pi t) \sin(2\pi 10^9 t)$$

$$\hat{x}(t_1, t_2) = \sin(2\pi t_1) \sin(2\pi 10^9 t_2)$$

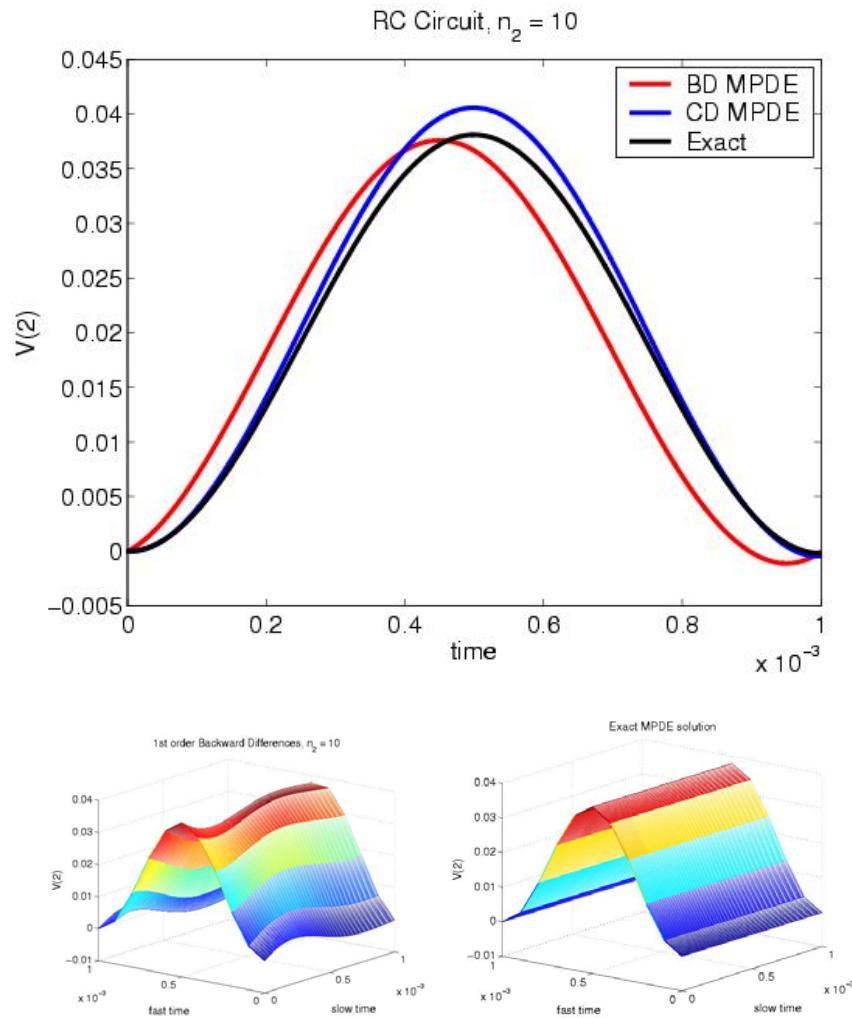
$$x(t) = \hat{x}(t, t)$$

MPDE Discretization

- Fast Time Discretization
Low Order & Coarse:**
 - high error
 - oscillation in slow time scale
- Mesh refinement study shows expected convergence**
- Win in convergence and speed by using higher order and/or greater refinement**

N2	Disc.	RelErr	Time(sec)
10	BD1	1.8e-1	46
	CD2	6.6e-2	28
	BD4	2.2e-2	27
40	CD4	4.6e-3	11
	BD1	4.6e-2	126
	CD2	3.8e-3	29
	BD4	1.2e-4	10
	CD4	5.7e-5	7

Myce Results, Todd Coffey



Epetra Communication

- Import/Export
 - Variable Block Communication **[Efficient Memory Usage]**
 - Efficient Buffering **[No Dynamic Memory]**
 - Direct Data Access **[No Search]**
- Impact
 - Huge reductions in buffer memory usage for key simulations
 - Critical Impact on Xyce Milestone Problem (Permafrost)

Epetra Imports	Memory (kB)	Time (s)	Solve Time (s)
Original	>3,000	~10	1.1
New	220	0.2	1.1

- Class of highly constrained problems now tractable for Salinas (C. Dohrmann)

EpetraExt(ensions)

- **Public Release: 3.0 & 4.0**
- **Capabilities**
 - **Transforms:** Singleton Filter, AMD, Remapping, Permutations
 - **Matrix Matrix:** Multiply, Add (Transpose)
 - A. Williams
 - **Block Manipulation:** Triangular Factorization, MPDE Support
 - **Distributed Boundary Resolution:** Generic Directories/Migrators
 - **Zoltan Interface:** Graph/Hypergraph Partitioning
 - **Graph Coloring:** Greedy, Lubi, DOF Ordering, Parallel
 - B. Spotz, R. Hooper
 - **Epetra Parallel I/O**
 - M. Heroux
- **Impact**
 - Xyce : Critical performance/robustness for ASC Level 1 Milestone
 - Premo, Charon: finite difference coloring
 - Zoltan: partitioning linear systems

Graph Coloring

Premo(Sierra) – Generated by Russel Hooper

		# colors	# dofs	Time (s)
B61 No-fins (Unstructured Tets)	Greedy, Largest First	155	33280	30.2
	Greedy, Smallest First	245	33280	26.5
	Greedy, Random	150	33280	25.2
	Lubi, Largest First	160	33280	24.9
	Lubi, Smallest First	153	33280	28.5
	Lubi, Random	156	33280	26.3
Sphere-Cone (Structured Hexes)	Greedy, Largest First	90	306820	69.3
	Greedy, Smallest First	105	306820	83.8
	Greedy, Random	84	306820	70.4
	Lubi, Largest First	84	306820	81.8
	Lubi, Smallest First	82	306820	85.0
	Lubi, Random	86	306820	86.8

New Capability Highlights

- Tim Davis's KLU in AMESOS (The “Clark Kent” of Direct Solvers)

- Gilbert/Peierel’s Left-Looking Sparse LU
- Fastest direct solver for Xyce circuits
- Block Triangular Factorization
 - » Based on our research ([D. Day](#))
- Available in next Xyce release
- Prototyping Dist. Mem. Impl.

Method	1 st Fac	Next Fac	LU nnz (10 ³)	MFlops
UmfPack	.440	.257	372	2.56
KSparse	84	14	294	0.90
KLU w/ BTF	.545	.02	201	0.16

- Zoltan Partitioning

- Weighted Graph Partitioning
 - » $\text{edgwt}(i,j) = F(|\text{val}[i][j]|)$
 - » Improved quality of block ILU
- Hypergraph Partitioning
 - » Improved model of communication cost
 - » Direct mapping to non-symmetric matrices
 - » Zoltan parallel algorithm in progress